Analysing the Impact of Learning Inputs - Application to Terrain Traversability Estimation

Ken Ho[1], Thierry Peynot[2] and Salah Sukkarieh[1]
[1] Australian Centre for Field Robotics (ACFR)
University of Sydney, NSW 2006, Australia
[2] Queensland University of Technology (QUT)
Brisbane QLD 4001, Australia

Abstract

Data-driven approaches such as Gaussian Process (GP) regression have been used extensively in recent robotics literature to achieve estimation by learning from experience. To ensure satisfactory performance, in most cases, multiple learning inputs are required. Intuitively, adding new inputs can often contribute to better estimation accuracy, however, it may come at the cost of a new sensor, larger training dataset and/or more complex learning, sometimes for limited benefits. Therefore, it is crucial to have a systematic procedure to determine the actual impact each input has on the estimation performance. To address this issue, in this paper we propose to analyse the impact of each input on the estimate using a variance-based sensitivity analysis method. We propose an approach built on Analysis of Variance (ANOVA) decomposition, which can characterise how the prediction changes as one or more of the input changes, and also quantify the prediction uncertainty as attributed from each of the inputs in the framework of dependent inputs. We apply the proposed approach to a terrain-traversability estimation method we proposed in prior work, which is based on multi-task GP regression, and we validate this implementation experimentally using a rover on a Mars-analogue terrain.

1 Introduction

With the increasing need to adapt to new environments, data-driven approaches such as Gaussian Process (GP) regression [Rasmussen and Williams, 2006] have been used extensively in recent robotics literature to achieve estimation and classification by learning from experience. In field robotics alone, examples of application include terrain modelling [Vasudevan *et al.*, 2009], occupancy mapping [O'Callaghan and Ramos, 2011] or terrain traversability estimation [Brooks and Iagnemma, 2012]. Multiple learning inputs are usually needed to adequately describe the various aspects that are correlated with the states that need to be estimated. These correlations are learnt using complex learning algorithms often considered as "black box" functions that provide little or no information about the actual impact of each input on the resulting estimate. Adding new inputs may be impactful but may come at a cost; for example, the robot may need to be fitted with additional sensors, learning is likely to be longer and more complex. Therefore, having a systematic procedure to determine the relevance of each learning input that is considered, by analysing the actual impact on the quality of the estimation, is essential. Without such procedure, it is difficult to understand the shortcomings of the system, or to evaluate the suitability of new inputs. Another benefit of such analysis would arise in case of the complete failure of a sensor on a robot, as we may be able to determine how much this may impact the quality of future estimates without having to train the system again with the new input configuration. This paper addresses this issue for regression techniques, in particular using GPs, by proposing to systematically analysing the impact of each input on the estimate using a variance-based sensitivity analysis method.

The overall accuracy of an estimator and the validity of the error can be checked using cross validation [Jones *et al.*, 1998]. However, this does not provide insight on the contribution or the impact of the learning inputs on the estimation. Sensitivity analysis methods can be used to better understand the responses of estimation systems [Saltelli *et al.*, 2000]. The Sobol index [Sobol, 1990], based on variance decomposition, measures sensitivity by expanding the global variance into partial variances. To validate the response of Gaussian Process (GP) regression, frameworks based on sensitivity analysis methods were proposed in [Schonlau and Welch, 2006]. These frameworks analysed the effects of input variables on the estimate. However, both of the above methods rely on

the assumption that the input variables are independent. If the input variables are dependent, the amount of response variance may be influenced by its dependence on other inputs, and thus can lead to incorrect interpretations [Mara and Tarantola, 2012].

To account for the contribution from dependent inputs, [Xu and Gertner, 2008] proposed to decompose the partial variance of an input into correlated and uncorrelated contribution components, assuming a linear effect from each component on the response. This approach was later extended by approximating the effect using a sum of functional components of low dimensions, and then computing the decomposition of response variance as a sum of partial variances [G. Li et al., 2010]. Sensitivity analysis methods to account for non-constant (heteroscedastic) variances in the estimate were proposed in [Dancik and Dorman, 2008].

In this paper we propose to analyse the impact of learning inputs using a sensitivity analysis method built on Analysis of Variance (ANOVA) decomposition in a framework of dependent inputs. This quantifies the contributions from each input to the variability of the resulting estimate, including the extent at which the estimate uncertainty can be attributed to the inputs. The method first decomposes the estimate into a multi-dimensional representation of primary and interaction effects between the inputs. The analytical sensitivity measure is then calculated for combinations of inputs, and indicates the significance of each input. This method has the potential to be used for any multi-input regression method where there is an uncertainty associated to the output. In this paper, we focus on GP-based regression.

We implement the proposed approach to analyse the impact of the learning inputs of a terrain traversability estimation method based on GPs, which was proposed in prior work. Terrain traversability estimation is critical for autonomous planetary rovers travelling on rough unstructured terrain, to provide them with the ability to anticipate situations that may compromise their safety and ability to conduct exploration missions. As the rover-terrain interaction in such terrain can be very difficult to model accurately, data-driven approaches have been developed to estimate terrain traversability by learning the rover's response on the terrain based on experience (e.g. [Brooks and Iagnemma, 2012]). Most recently, in [Ho *et al.*, 2013a] the authors proposed a new approach to predict the angles of attitude and chassis configuration (i.e. the joint angles) of a rover on any location of deformable terrain, given incomplete terrain elevation data. Predicting these angles is essential to traversability estimation, as they reflect the difficulty that the rover may have to traverse the terrain. On very rough terrain, this prediction can even allow us to anticipate risks for the stability of the platform.

To achieve this prediction on specific locations of the terrain, many inputs are potentially useful to learn from. The inputs used in the approach presented in [Ho *et al.*, 2013a] include the attitude and configuration angles from experience and the local curvatures of predicted angles, which give an image of the local variation of terrain geometry (we showed in [Ho *et al.*, 2013a] that this information was useful to anticipate potential terrain deformation and better estimate the resulting rover attitude and chassis configuration). However, other inputs such as the velocity of the rover, may also be relevant in the learning process. In this paper, we apply the proposed sensitivity analysis method to this data-driven estimation of terrain traversability with inclusion of the velocity of the rover, and we validate it experimentally with data collected using a prototype rover on a Mars-analogue terrain. We quantify the impact that the inputs of the original method have on the estimate, as well as the additional input considered in this study, thereby validating the choice to add this new input. To the best of our knowledge, this is the first time such sensitivity analysis is conducted on a learning-based terrain traversability estimation method.

The paper is organised as follows. Sec. 2 formulates the proposed sensitivity analysis method. Sec. 3 outlines the background learning framework we used as application: a method to predict a rover's attitude and chassis configuration on unstructured terrain. Sec. 4 describes the implementation of the approach and our experimental setup. In Sec. 5 we present and discuss the experimental results obtained. Finally, Sec. 6 concludes the paper.

2 Proposed Sensitivity Analysis Approach

2.1 Analysis of Variance (ANOVA) decomposition

We employ a functional ANOVA approach to breakdown the estimate into a linear combination of contributions from the learning inputs. The percentage of total contribution attributed among the inputs then provides a measure of importance of the interaction effect between each input and the resulting estimate [Dancik and Dorman, 2008]. This quantifies the contributions from each input to the variability of the resulting estimate, including the extent at which the estimate uncertainty can be attributed to the inputs. In our application, ANOVA decomposes the total mean and variance of the GP estimator into contributions from dependent inputs.

The concept of the ANOVA decomposition is illustrated with an example in Fig. 1. Consider we have three learning inputs that contribute to the total estimate (see Fig. 1(a)), this approach can tell us that In-

(a) Primary effects
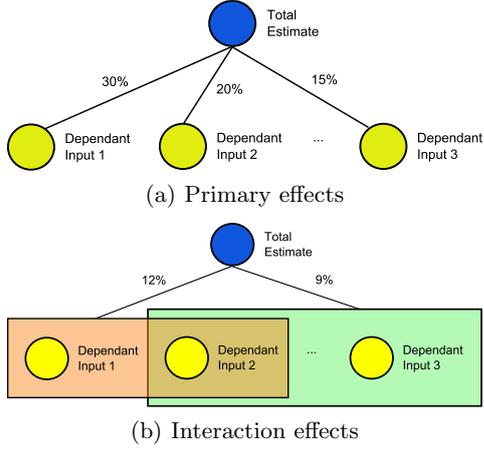

(b) Interaction effects

Figure 1: ANOVA approach to breakdown the estimate into a linear combination of contributions from learning inputs. The orange and green boxes indicate the interaction effects between Inputs 1 and 2, and between Inputs 2 and 3, respectively.

put 1 alone contributed 30% to the total estimate, Input 2 contributed 20%, and Input 3 contributed 15%. Expanding the idea further, we also want to account for the interaction effects between dependent learning inputs. That is, we want to see the effects of a subset of the inputs, and also how they interact with each other to contribute to the total estimate. For example, if we consider the interaction effects between Inputs 1 and 2, and between Inputs 2 and 3 (see Fig. 1(b)), this approach can tell us that the first pair contributed to 12% of the total estimate, and the second pair contributed to 9%.

Let us consider a d-dimensional vector of input variables $\boldsymbol{x} = (x_1, ..., x_d)^T$ $(d > 2)$. The multi-input single-output regression problem is to find an estimate of the output variable $y(\boldsymbol{x})$ for any value of $\boldsymbol{x}$, in particular values that are not available in the training data. With several output variables, each is treated separately. Suppose that we are interested in the effect of a subset of input variables, held in a vector $\boldsymbol{x}_e$, where e denotes the set of subscripts of the variables of interest. The vector of remaining variables is denoted by $\boldsymbol{x}_{-e}$, such that $\boldsymbol{x} = (\boldsymbol{x}_e, \boldsymbol{x}_{-e})$ [Schonlau and Welch, 2006]. For example, if our interest lies in the effects of x_1 and x_2, we have $\boldsymbol{x}_e = (x_1, x_2)$, while $\boldsymbol{x}_{-e} = (x_3, ..., x_d)$. Generalising this notation, $x_e = (x_{i_1}, ..., x_{i_n})$ where $(i_1, ..., i_n)$ are n indices taken among d, and x_e represents n components of $\boldsymbol{x}$. Let us define S as all non-empty subsets of $\{1, ..., d\}$. $\{\boldsymbol{x}_e\}_{e \in S}$ represents all possible combinations of input variables.

We can then consider the output estimate as an inte-

grable function $y(\boldsymbol{x})$, thus expressed as:

$$y(\boldsymbol{x}) = y_0 + \sum_{e \in S} y_e(\boldsymbol{x}_e). \tag{1}$$

where $y_e(\boldsymbol{x}_e)$ is a component of the estimate output y that is dependent on the subset of input variables of interest $\boldsymbol{x}_e$. We may then define an effect from a learning input by "integrating out" the other inputs. Under certain conditions, this leads to a simple decomposition of $y(\boldsymbol{x})$ into contributions from various effects, with a corresponding decomposition of the total variance of $y(\boldsymbol{x})$ over an input region of interest $\mathcal{X}$ [Schonlau and Welch, 2006]. Moreover, these effects and their variance contributions can be easily estimated.

To decompose the resulting estimate, first we need to find the marginal effect $\bar{y}_e(\boldsymbol{x_e})$. This is the overall effect of all variables in $\boldsymbol{x}_e = (x_{i_1}, ..., x_{i_n})$ on the estimate, and is defined by integrating out all other variables [Schonlau and Welch, 2006]:

$$\bar{y}_e(\boldsymbol{x}_e) = \int_{\otimes_{j \notin e} \mathcal{X}_j} y(\boldsymbol{x}_e, \boldsymbol{x}_{-e}) \prod_{j \notin e} w_j(x_j) dx_j \text{ for } \boldsymbol{x}_e \in \otimes_{j \in e} \mathcal{X}_j,$$
$$\tag{2}$$

where $w_j(x_j)$ is a weight function that represents the relative interest for variable x_j among the components of $\boldsymbol{x}$, $\mathcal{X}_j$ denotes the values of interest for variable x_j and $\otimes_{j \in e} \mathcal{X}_j$ is a direct product of one-dimensional regions. To represent uniform interest across the range of values for x_j, in this paper we choose the values of $w_j(x_j)$ to be equal.

We then use Eq. (2) to decompose the resulting estimate $y(\boldsymbol{x})$ into corrected effects involving the contributions from any number of variables in $\boldsymbol{x} \in \mathcal{X}$:

$$y(\boldsymbol{x}) = \mu_0 + \sum_{j=1}^{d} \mu_j(x_j) + \sum_{j=1}^{d-1} \sum_{j'=j+1}^{d} \mu_{jj'}(x_j, x_{j'}) + \dots$$
$$+ \mu_{1\dots d}(x_1, \dots, x_d), \tag{3}$$

where μ_0 is the overall average, $\mu_j(x_j)$ is the corrected primary effect of a single variable x_j, $\mu_{jj'}(x_j, x_{j'})$ is the corrected interaction effect of two variables x_j and $x_{j'}$ $(j \neq j')$, and so on [Schonlau and Welch, 2006]:

$$\mu_0 = \int_{\mathcal{X}} y(x)w(x)dx$$
$$\mu_j(x_j) = \bar{y}_j(x_j) - \mu_0 \qquad \text{for } x_j \in \mathcal{X}_j$$
$$\mu_{jj'}(x_j, x_{j'}) = \bar{y}_{jj'}(x_j, x_{j'}) - \mu_j(x_j) - \mu_{j'}(x_{j'}) - \mu_0$$
$$\text{for } x_j, x_{j'} \in \mathcal{X}_j \otimes \mathcal{X}_{j'} \tag{4}$$

For example, to examine the contributions from inputs x_1 and x_2 on the resulting estimate, we can consider the

overall joint effect:

$$\bar{y}_{12}(x_1, x_2) = \mu_0 + \mu_1(x_1) + \mu_2(x_2) + \mu_{12}(x_1, x_2)$$
$$\text{for } x_1, x_2 \in \mathcal{X}_1 \otimes \mathcal{X}_2 \quad (5)$$

which is a linear combination of the overall average (mean) of the estimation, the primary effect of inputs x_1 and x_2 individually, and the secondary effects of the interaction between these two inputs.

In practice, we first estimate the marginal effects $\bar{y}_e(\boldsymbol{x}_e)$ using a best linear unbiased predictor (BLUP) [Schonlau and Welch, 2006], and then compute the corresponding estimated corrected effect by subtracting all estimated lower-order corrected effects. Using this decomposition, we can determine the impact of the learning inputs on the resulting estimate as a function of its interaction with other inputs.

2.2 Numerical Evaluation of the Marginal Effects

Following the approach in [Schonlau and Welch, 2006], $y(\boldsymbol{x})$ can be treated as a realisation of the random function $Y(\boldsymbol{x})$:

$$Y(\boldsymbol{x}) = \boldsymbol{f}'(\boldsymbol{x})\beta + Z(\boldsymbol{x}) \quad (6)$$

where $\boldsymbol{f}(\boldsymbol{x}) = [f_1(\boldsymbol{x}), ..., f_h(\boldsymbol{x})]'$ is a vector of h known regression functions, β is a $h \times 1$ vector of parameters to be estimated, and Z is a Gaussian stochastic process indexed by $\boldsymbol{x}$. It then follows that $\bar{y}_e(\boldsymbol{x}_e)$ is a realisation of the analogously integrated random function $\bar{Y}_e(\boldsymbol{x}_e)$:

$$\bar{Y}_e(\boldsymbol{x}_e) = \bar{f}'_e(\boldsymbol{x}_e)\beta + \bar{Z}_e(\boldsymbol{x}_e) \qquad \text{for } \boldsymbol{x}_e \in \otimes_{j\in e}\mathcal{X}_j \quad (7)$$

where $\bar{f}_e(\boldsymbol{x}_e)$ and $\bar{Z}_e(\boldsymbol{x}_e)$ can be expressed as integrals as in Eq. (2):

$$\bar{f}_e(\boldsymbol{x}_e) = \int_{\otimes_{j\notin e}\mathcal{X}_j} f(\boldsymbol{x}_e, \boldsymbol{x}_{-e}) \prod_{j\notin e} w_j(x_j)dx_j \text{ for } \boldsymbol{x}_e \in \otimes_{j\in e}\mathcal{X}_j,$$
$$(8)$$

$$\bar{Z}_e(\boldsymbol{x}_e) = \int_{\otimes_{j\notin e}\mathcal{X}_j} Z(\boldsymbol{x}_e, \boldsymbol{x}_{-e}) \prod_{j\notin e} w_j(x_j)dx_j \text{ for } \boldsymbol{x}_e \in \otimes_{j\in e}\mathcal{X}_j.$$
$$(9)$$

We then solve for f_e and Z_e in Eqs. (8) and (9) using BLUP [Schonlau and Welch, 2006]. From the estimated marginal effects, the primary and interaction effects can then be estimated using Eq. (4).

2.3 Analytical Sensitivity Measure (ASM)

Global sensitivity indices are used for estimating the influence of individual or groups of variables on the resulting estimate. A common example of a global sensitivity index is the Sobol index [Sobol, 1990]. However, it does not quantify the uncertainty brought by dependent inputs. From the decomposition of the mean of the

estimate, an Analytical Sensitivity Measure (ASM) was proposed in [Durrande *et al.*, 2013] to quantify the contribution of each *independent* learning input x_j to the estimate $y(\boldsymbol{x})$:

$$S_j = \frac{V\left[\mu_j(x_j)\right]}{V\left[y(\boldsymbol{x})\right]} \quad (10)$$

For *dependent* input variables, the ASM S_j that accounts for both the estimate mean and uncertainty can be computed as [Chastaing and Gratiet, 2013]:

$$S_j = \frac{V[\mu_j(x_j)] + Cov[\mu_j(x_j), \mu_{j^c}(\boldsymbol{x})]}{V[y(\boldsymbol{x})]}, \quad (11)$$

where $\mu_{j^c}(\boldsymbol{x}) = y(\boldsymbol{x}) - \mu_j(x_j)$. In our implementation, we compute S_j for each learning input to determine the impact of each x_j on the resulting estimate. Note that an ASM can also be derived for the interaction effects.

3 Application: Rover Attitude Prediction in Partially Occluded and Deformable terrain

Recent literature showed that an effective way to predict the rover's response on upcoming terrain is to learn the correlation between exteroceptive and proprioceptive sensor information. This concept is known as *near-to-far* learning and was demonstrated in various references, such as [Brooks and Iagnemma, 2012; Krebs *et al.*, 2010]. In [Ho *et al.*, 2013a] the authors proposed a near-to-far approach to predict the attitude and configuration angles of a rover while addressing the problems of incomplete terrain data and of terrain deformation, by using two connected components (see Fig. 2). Given incomplete terrain data, the first component, named *Rigid-Terrain Traversability Estimation* (R-TTE), provides an initial estimate of the rover configuration and attitude, Φ^*_{rigid}, *before* any terrain deformation may occur [Ho *et al.*, 2013b]. This is equivalent to assuming that the terrain is rigid. The second component, *Rigid-to-Deformable Terrain Traversability Estimate* (R2D-TTE), then refines this prediction by accounting for the effects of terrain deformation on rover configuration and attitude. The final estimate is Φ^*_{deform}. Both processes are stochastic. For convenience, in the remainder of the paper, rover configuration will refer to *both* attitude and chassis configuration angles.

3.1 R-TTE

The R-TTE module addresses the problem of incomplete terrain data [Ho *et al.*, 2013b]. It estimates a complete map of Φ^*_{rigid} by performing GP regression over an incomplete map. This approach exploits the explicit correlation in rover configuration during operation by learning a kernel function from experience. The estimation
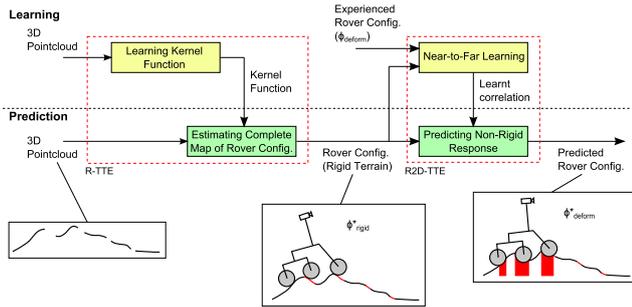
Figure 2: Our approach to predict rover attitude and configuration on terrain. Given an incomplete point cloud, R-TTE makes an initial estimate of the configuration over the entire map, assuming the terrain is rigid. R2D-TTE then refines this estimate to account for possible terrain deformation.

scenario is set up as a GP regression problem to predict $\Phi^*_{rigid}(x, y, \psi)$ at each position (x, y) on a Digital Elevation Map (DEM) over different heading angles ψ. The GP posterior (estimate) $y = f^*$, represented by the mean $\bar{f}^*$ and covariance $cov(f^*)$ can be given as:

$$\bar{f}^* = K(\boldsymbol{x}^*, \boldsymbol{x})[K(\boldsymbol{x}, \boldsymbol{x}) + \sigma_n^2 I]^{-1} z$$

$$cov(f^*)$$
$$= K(\boldsymbol{x}^*, \boldsymbol{x}^*) - K(\boldsymbol{x}^*, \boldsymbol{x})[K(\boldsymbol{x}, \boldsymbol{x}) + \sigma_n^2 I]^{-1} K(\boldsymbol{x}, \boldsymbol{x}^*)$$

where K represents the covariance matrix evaluated using the learnt kernel function at all pairs of training points $\boldsymbol{x}$ and query points $\boldsymbol{x}^*$, σ_n is the noise variance, and z is the training target.

3.2 R2D-TTE

The R2D-TTE module, introduced in [Ho *et al.*, 2013a], refines the estimate provided by R-TTE by accounting for the effects of terrain deformation. In this paper we extended the estimation process to exploit the local variations in Φ_{rigid} that correlate with the actual rover configuration resulting from terrain deformation, i.e. Φ_{deform}, and included driving speed as an additional learning input. This idea is implemented by learning the correlation between the initial prediction, Φ^*_{rigid}, its local variations, and experience in Φ_{deform} collected during training (see Fig. 3). During learning, the rover observes a patch of terrain and predicts Φ^*_{rigid}. When the rover traverses over this terrain, it learns the correlation between Φ^*_{rigid} and the experienced rover configuration Φ_{deform} with terrain deformation. Once the training is complete, in operation, the rover uses the learnt correlations to predict Φ^*_{deform} from new exteroceptive data.

Learning is performed in a multi-task heteroscedastic GP framework that considers the interaction between
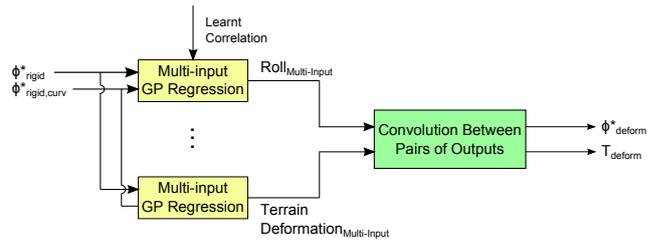

Figure 3: R2D-TTE process to account for the effects of terrain deformation on rover configuration, using correlations learnt in experiments.

multiple training inputs and targets. We use multi-input GP regression by Automatic Relevance Determination (ARD) to learn the correlation between the training inputs $\boldsymbol{x}$ and each component in the target z. We use convolution processes to account for the correlations between estimation outputs [Caruana, 1997]. This approach uses a convolution between a smoothing kernel k_q and latent functions $u(z)$ to express each output f_q:

$$f_q(\boldsymbol{x}) = \int_{-\infty}^{\infty} k_q(\boldsymbol{x} - z) u(z) dz$$

We then use multiplication of Gaussian distributions to determine the correlation between pairs of outputs as well as between any given output and the latent function. Using these covariance matrices we perform joint-prediction of the estimation outputs. For more details please refer to [Ho *et al.*, 2013a].

4 Implementation and Experimental Setup

We implemented the ANOVA decomposition to analyse the estimation results from R2D-TTE. Training and validation data were obtained with a prototype planetary rover navigating on a Mars-analogue terrain.

4.1 GP Learning Inputs, Targets and Outputs

The training input $\boldsymbol{x}$ used in the GP regression within R2D-TTE includes $\Phi^*_{rigid} = \{\phi, \theta, \alpha_1, \alpha_2\}$, where ϕ and θ are the roll and pitch of the rover, respectively, and the α_i are angles of the joints of the chassis (see Fig. 4(b)). It also includes the local curvatures of Φ^*_{rigid}, i.e. $\Phi^*_{rigid,curv}$, defined as the combined planform and platform curvature of the rover configuration on the terrain (see details [Ho *et al.*, 2013a]). In this paper we consider an additional input, the rover driving speed v_{rover}, to analyse its potential impact on the estimate:

$$\boldsymbol{x} = [\phi, \phi_{curv}, \theta, \theta_{curv}, \alpha_1, \alpha_{1_{curv}}, \alpha_2, \alpha_{2_{curv}}, v_{rover}].$$

The training target $z = [\Phi_{deform}, \mathcal{T}_{deform}]$ includes the actual rover configuration Φ_{deform} and the curvature of Φ on *deformed* terrain, $\mathcal{T}_{deform}$, which conveys information on the effect of terrain deformation on configuration curvature.

$$
\begin{aligned}
\Phi_{deform} &= \left[\phi_{deform}, \theta_{deform}, \alpha_{1_{deform}} \alpha_{2_{deform}} \right], \\
\mathcal{T}_{deform} &= \left[\phi_{curv}, \theta_{curv}, \alpha_{1_{curv}}, \alpha_{2_{curv}} \right]_{deform},
\end{aligned}
$$

In this paper, we focus on the analysis of the impact of the different inputs in $\boldsymbol{x}$ on the estimation of Φ_{deform}.

4.2 Platform - *Mawson* Rover

The experiments were conducted using *Mawson*, a 6-wheeled prototype rover platform with a rocker-bogie chassis (see Fig. 4(a)). Mawson is approximately $80cm$ long, $63cm$ wide, and $90cm$ tall. The radius of each wheel is $5cm$. It is equipped with the following sensors. Two colour cameras and a RGB-D camera mounted on a pan-tilt unit, tilted down $\approx 20°$. The acquired 3D point clouds are used to generate DEMs. Two Hall-effect encoders measure the two rear bogie angles (α_1, α_2 in Fig. 4(b)). An Intersense^{TM} IS-1200 fuses data from a visual camera and an inertial measurement unit to provide the 6-DOF sensor pose, including rover attitude, with an average accuracy of $2cm$ in position and $1°$ in orientation/attitude.

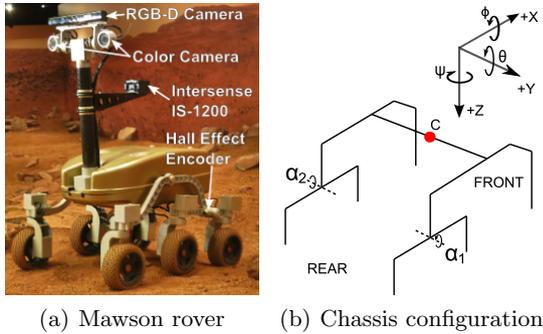| (a) Mawson rover | (b) Chassis configuration |

Figure 4: Experimental rover platform.

4.3 Test Environment

We conducted our experiments at the Marsyard, a Mars-analogue terrain in Sydney, Australia (see Fig. 5). The Marsyard is approximately $15m \times 8m$ and contains slopes, soil and rocks similar to Martian terrain. The typical obstacle size in the Marsyard is approximately $0.05m$ to $0.2m$ in radius. Combined with the mixed sizes in gravel granules, this represents a considerable challenge in traversability for Mawson since its wheel radius is $0.05m$.

Figure 5: Section of the Marsyard in Sydney, Australia.

4.4 Experimental Data For Learning

We performed a range of traversals over different terrain to engage Mawson in a variety of situations that it is likely to encounter during operation. Before the rover traversed on the terrain, we recorded the point cloud of the terrain using an external depth sensor. As the rover traversed the terrain, we collected the experienced rover configuration Φ_{deform} using the Intersense sensor, as well as terrain data using the onboard depth sensor. After terrain traversal, we acquired another point cloud of the terrain using the external depth sensor. To quantify terrain deformation, we compared the DEMs generated from terrain data acquired *before* and *after* rover traversal, after aligning them using Nearest-Neighbour Iterative Closest-Point.

5 Experimental Results

We used R2D-TTE to predict rover configuration Φ_{deform}^* with terrain deformation. We then evaluated the ability of the ANOVA approach to quantify the impact of the learning inputs on the estimate. This was performed in two steps. We first decomposed the rover configuration estimate into a sum of contributions from each input. We then evaluated the ASM for each input as an overall measure of the input's influence on the rover configuration estimate.

The rover was driven at different speeds ($0.5m/s$ to $1m/s$) on varying terrain geometry and nature, to observe the impact of v_{rover}. We also investigated the impact that interaction effects between the inputs have on the estimate, and how they vary as the experimental situation changes. The ground truth of rover configuration was obtained from the Intersense sensor (rover attitude) and Hall-effect encoders (chassis configuration), see Sec. 4.2. During the experiments, the driving speed was also recorded, and the ground truth of terrain deformation was obtained using the external depth sensor, as described in Sec. 4.4. In this validation process, the GP hyperparameters were trained with 2500 training points. This was chosen based on an experimental analysis of the learning rate of the algorithm. The cross-validation was

then performed using a distinct set of 3000 points, which correspond to about $22m$ of terrain traversal (about 28 times the length of the rover).

Fig. 6 illustrates the rover attitude estimated using R-TTE and R2D-TTE in a section of traversal, with the rover operating at different speeds. The estimation made by R2D-TTE is able to anticipate the effects of terrain deformation. This is particularly clear between 12.5 and $19m$ distance travelled (section 'A' in Fig. 6), where there was significant terrain deformation (see the large errors made by R-TTE in this segment). However, it is unclear which of the inputs are contributing positively to the learning process, and which ones are leading the estimate astray. For example, between 10 and $12m$ distance travelled (section B in Fig. 6), it is unclear why the estimate of rover roll deviates from the ground truth that remained fairly constant. On the other hand, the estimates of rover roll and pitch between 14 and $17m$ (section C) are less erroneous than between 10 and $12m$, albeit in an area with a larger range of roll and pitch variations.

5.1 Decomposing the Rover Configuration Estimate Φ^*_{deform}

We first investigate the primary effects ($\mu_j(x_j)$) from each learning input x_j. Fig. 7 shows a decomposition of the primary effects with the highest impact on the resulting estimate (see Table 1). We can identify sections along the rover's traversal in Fig. 7 where each input is the main contributor to the estimate. In particular, in Fig. 7(b), we can identify that the components of Φ^*_{rigid}, i.e. $\{\phi_{rigid}, \theta_{rigid}, \alpha_{1rigid}, \alpha_{2rigid}\}$, contribute significantly to the overall estimate (black line) between $14m$ and $17m$ distance travelled. From Figs. 7(a) and (b), we can also see that the primary contribution from Φ^*_{rigid} has an adverse effect on the overall estimate of ϕ_{deform} and θ_{deform} between 12 and $16m$, where there was significant terrain deformation. This indicates that other interaction effects between the dependent inputs are contributing to the ability of the R2D-TTE approach to anticipate terrain deformation.

We then investigate the interaction effects from the learning inputs. From the analysis of ASM of each input in Sec. 5.2, we selected the sets of inputs with the three sets of highest pair-wise and triplet-wise ASM values.

Fig. 8 shows a decomposition of the 2^{nd} order interaction effects with the highest impact on the resulting estimate of rover pitch (see Table 2). It can be seen that the dominant interaction effect changes among the combination of inputs along the rover's trajectory. For example, $(\phi, \theta)_{rigid}$, $(\theta, \alpha_1)_{rigid}$, and $(\alpha_1, \alpha_2)_{rigid}$ are the dominant 2^{nd} order interaction effects from $16m$ to $20m$ for the Pitch θ. Also, comparing the overall estimate of ϕ_{deform} and θ_{deform} in Fig. 8 with Fig. 7(b), we can

see that the 2^{nd} order interaction effects are much closer to the overall effects from all inputs in x (black line in Fig. 8). This suggests that the 2^{nd} order interaction effects are much more effective than the primary effects from the dependent inputs at anticipating terrain deformation between 12 and $16m$ in this test.

We also analysed the 3^{rd} order effects (the figure is not shown due to lack of space). The analysis showed that the dominant interaction effect also changed among the combination of inputs along the rover's trajectory. For example, $((\phi, \theta)_{rigid}, v_{rover})$ was the dominant 3^{rd} order interaction effect from $12m$ to $14m$ in the test shown in Fig. 6, while $(\phi, \theta, \phi_{curv})_{rigid}$ and $(\phi, \alpha_1, \phi_{curv})_{rigid}$ were the dominant 3^{rd} order effects from $15m$ to $20m$.

By considering the decomposition of the rover configuration estimate Φ^*_{deform} into contributions from each input, we can identify regions where each input is contributing positively or negatively to the learning process. As a result, we are able to systematically analyse the effects of each existing input, as well as additional inputs, on the estimate.

5.2 Analytical Sensitivity Measure for Each Learning Input

We determined the ASM S_j that accounts for the mean μ_j and uncertainty $V(\mu_j)$ of the decomposed contributions from each learning input to quantify their contribution to the resulting estimate. Table 1 shows the analytical sensitivity measure of primary effects from each input to Φ_{deform}. It can be seen that Φ_{rigid} in the inputs contributes to the highest values in first order effects in Φ_{deform} (highlighted in Table 1). This is because Φ_{rigid} is expected to be very similar to Φ_{deform} in areas with minimal terrain deformation. However, the contribution of v_{rover} is also significant, which validates the choice of adding driving speed as a learning input, compared to the original method in [Ho *et al.*, 2013a].

Table 1: Analytical sensitivity measure S_j of primary (first order) effects from each learning input to Φ_{deform}.

	ϕ_{deform} (%)	θ_{deform} (%)	$\alpha_{1deform}$ (%)	$\alpha_{2deform}$ (%)
ϕ_{rigid}	**13.4**	7.7	3.9	4.1
θ_{rigid}	9.8	**12.5**	9.2	7.8
α_{1rigid}	4.3	0.9	**7.1**	3.7
α_{2rigid}	3.4	4.8	2.8	**8.2**
$\phi_{rigid,curv}$	3.7	5.1	1.1	1.8
$\theta_{rigid,curv}$	2.1	0.8	1.9	0.7
$\alpha_{1rigid,curv}$	1.2	2.7	4.8	1.9
$\alpha_{2rigid,curv}$	2.9	1.1	4.2	3.8
v_{rover}	5.1	4.2	2.3	3.4
$\sum primary$	45.9	39.8	37.3	35.4

Table 2 shows the analytical sensitivity measure of se-

(a) Roll (ϕ)
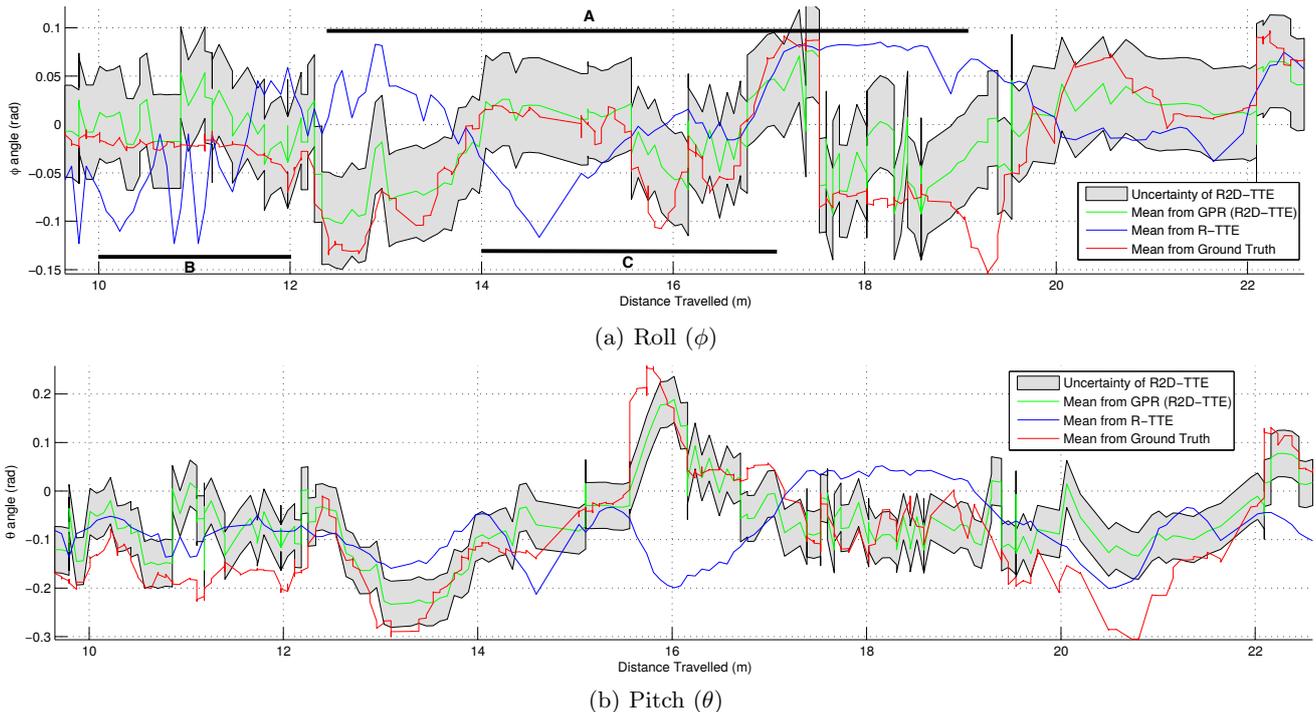

(b) Pitch (θ)

Figure 6: Prediction of each rover attitude angle using R2D-TTE (mean in green, with uncertainty in grey), compared with ground truth (in red) and the R-TTE estimate (in blue) with rigid terrain assumption (due to lack of space, only the attitude angles are shown).

lected interaction effects from combinations of inputs to Φ_{deform} (for conciseness we only show rows that contain elements higher than 2.5%). The interaction effects with the highest impact on the estimated rover roll and pitch are highlighted, and the sum of the selected interaction effects is shown at the bottom of the table. It can be seen that the interaction effects of v_{rover} with α_1 and α_2 are among the most impactful. The interaction effects of $\Phi_{rigid,curv.}$ with other inputs are also significant, having an analytical sensitivity measure between 50 to 65% of the highest values in the estimate of ϕ_{deform} and θ_{deform}. It should be noted that other combinations of interaction effects also contribute to the resulting estimate Φ_{deform}, but are minor and thus not shown in Table 2 for clarity.

6 Conclusion

In this paper, we proposed a sensitivity analysis approach built on Analysis of Variance (ANOVA) decomposition to analyse the impact of different learning inputs on the estimate produced by a multi-input regression technique used for estimation, such as Gaussian process regression. The method first decomposes the resultant estimate into a multi-dimensional representation of primary and interaction effects between the inputs, and

Table 2: Analytical sensitivity measure of selected 2^{nd} order interaction effects from each learning input to Φ_{deform}.

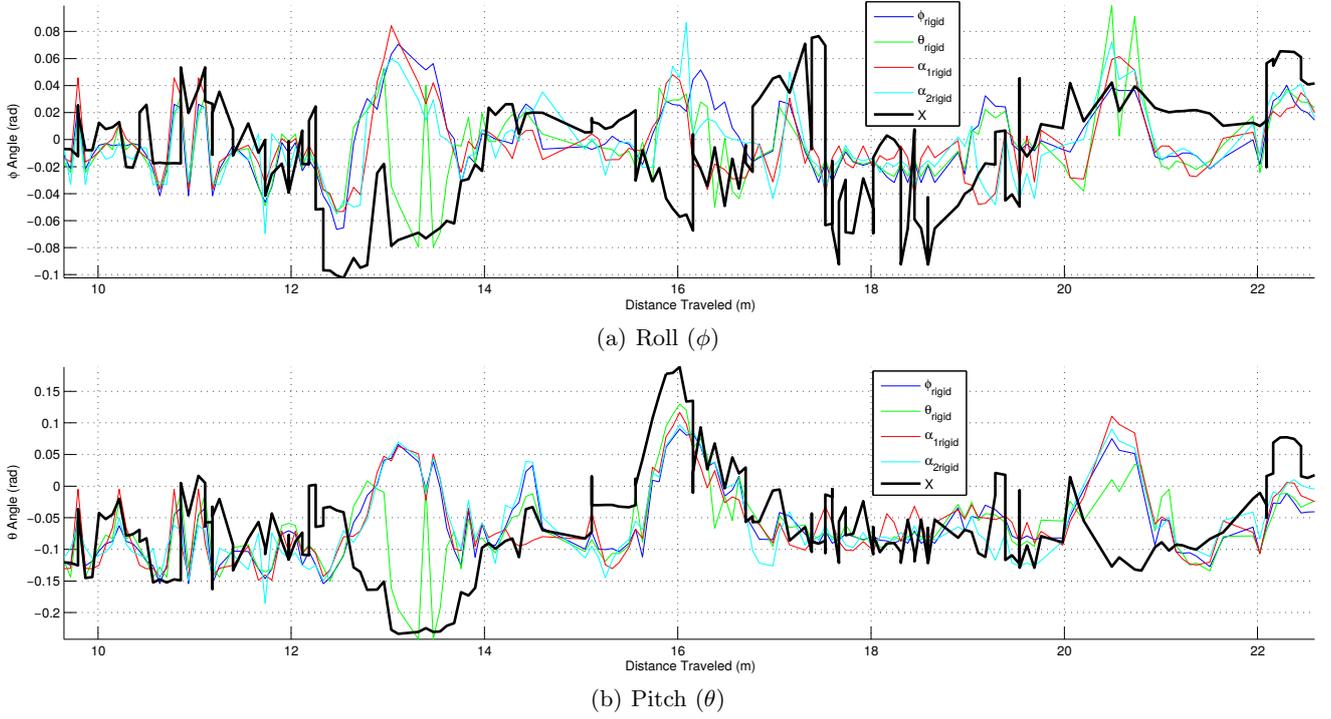|	$\phi_{def.}$ (%)	$\theta_{def.}$ (%)	$\alpha_{1_{def.}}$ (%)	$\alpha_{2_{def.}}$ (%)
$(\phi, \theta)_{rigid}$	**5.1**	**4.4**	**4.3**	**4.1**
$(\phi, \alpha_1)_{rigid}$	2.8	**4.3**	3.7	2.9
$(\phi, \alpha_2)_{rigid}$	**2.9**	2.3	3.6	2.1
$(\theta, \alpha_1)_{rigid}$	1.9	1.3	**4.2**	**4.7**
$(\theta, \alpha_2)_{rigid}$	**3.1**	2.1	**5.3**	**5.9**
$(\theta, \phi_{curv})_{rigid}$	0.2	3.3	0.2	0.1
$(\theta, \theta_{curv})_{rigid}$	0.2	3.8	0.1	0.1
$(\alpha_1, \alpha_2)_{rigid}$	0.5	0.3	**4.2**	3.2
$(\alpha_1, \alpha_{1,curv})_{rigid}$	1.5	2.1	**4.2**	0.2
$(\alpha_1)_{rigid}, v_{rover}$	0.6	3.7	3.2	0.1
$(\alpha_2, \theta_{curv})_{rigid}$	0.2	0.1	0.1	3.4
$(\alpha_2, \alpha_{2,curv})_{rigid}$	1.6	1.6	0.1	**4.1**
$(\alpha_2)_{rigid}, v_{rover}$	0.5	**3.9**	1.4	0.1
$\sum interaction, x \subset \mathcal{X}$	31	40.9	41	38.7

(a) Roll (ϕ)

(b) Pitch (θ)

Figure 7: Decomposition of primary effects on the prediction of rover attitude angles. The black line shows the overall effects from *all* learning inputs in $\boldsymbol{x}$. The coloured lines represent the primary effects from individual inputs (blue: $\mu(\phi_{rigid})$, green: $\mu(\theta_{rigid})$, red: $\mu(\alpha_{1_{rigid}})$, aqua: $\mu(\alpha_{2_{rigid}})$). The closer each coloured line (primary effect) is from the black line (overall effects), the more influential the corresponding input is.
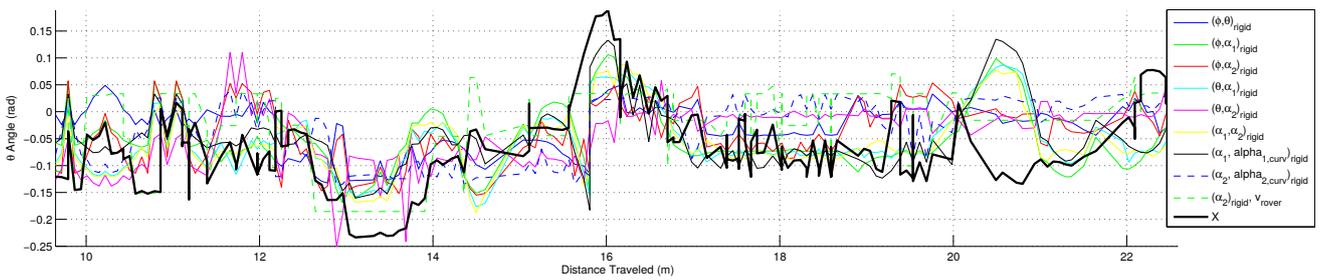

Figure 8: Decomposition of 2^{nd} order interaction effects $\mu_{jj'}(x_j, x_{j'})$ on the prediction of the Pitch of the rover (θ). The overall effects from *all* learning inputs in $\boldsymbol{x}$ are shown as the black line and each coloured line represents a 2^{nd} order interaction effect. The closer a coloured line is from the black line, the more substantial the corresponding 2^{nd} order interaction effects from that pair of inputs is.

then calculates the analytical sensitivity measure that indicates the significance of each input. Since adding extra inputs to a learning framework may come at a significant cost, such analysis allows us to select the most useful inputs only. We demonstrated this approach on a technique to predict a rover's attitude and chassis configuration on unstructured terrain from incomplete terrain data, while accounting for the effects of terrain deformation, which was introduced by the authors in prior work.

We showed that terrain geometry, reflected by the rover's attitude and configuration, was the most informative input data, having the highest analytical sensitivity measures as primary effect. It also had significant impact as an interaction effect, when combined with driving speed. We validated the addition of the rover speed as an extra input, by quantifying the significance of the impact it had on the estimate.

In future work, this analytical approach will allow us to systematically consider the utility of other potential learning inputs we are considering, such as terrain colour and texture. We will also be able to determine other types of possible shortcomings of the learning approach by drawing a connection between contributions from each input and physical and/or spatial phenomena such as terrain deformation.

Acknowledgment

References

[Brooks and Iagnemma, 2012] C. A. Brooks and K. Iagnemma. Self-supervised terrain classification for planetary surface exploration rovers. *J. Field Robotics*, 29(3):445–468, 2012.

[Caruana, 1997] R. Caruana. Multitask learning. *Machine learning*, 28(1):41–75, 1997.

[Chastaing and Gratiet, 2013] G. Chastaing and L. L. Gratiet. ANOVA decomposition of conditional gaussian processes for sensitivity analysis with dependent inputs. *arXiv:1310.3578*, 2013.

[Dancik and Dorman, 2008] G. M. Dancik and K. S. Dorman. Statistical analysis for computer models of biological systems using R. *Bioinformatics*, pages 1966–1967, 2008.

[Durrande et al., 2013] N. Durrande, D. Ginsbourger, O. Roustant, and L. Carraro. Reproducing kernels for spaces of zero mean functions. application to sensitivity analysis. *J. Multivariate Analysis*, 115:57–67, 2013.

[G. Li et al., 2010] G. Li et al. Global sensitivity analysis for systems with independent and/or correlated inputs. *J. Physical Chemistry A*, 114(19):6022–6032, 2010.

[Ho et al., 2013a] K. Ho, T. Peynot, and S. Sukkarieh. A near-to-far non-parametric learning approach for estimating traversability in deformable terrain. In *IEEE/RSJ Int. Conf. on Robotics and Intelligent Systems*, 2013.

[Ho et al., 2013b] K. Ho, T. Peynot, and S. Sukkarieh. Traversability estimation for a planetary rover via experimental kernel learning in a gaussian process framework. In *IEEE Int. Conf. on Robotics and Automation*, 2013.

[Jones et al., 1998] D. R. Jones, M. Schonlau, and W. J. Welch. Efficient global optimization of expensive black-box functions. *J. Global optimization*, 13(4):455–492, 1998.

[Krebs et al., 2010] A. Krebs, C. Pradalier, and R. Siegwart. Adaptive rover behaviour based on online empirical evaluation: Rover-terrain interaction and near-to-far learning. *J. Field Robotics*, 27(2), 2010.

[Mara and Tarantola, 2012] T. A. Mara and S. Tarantola. Variance-based sensitivity indices for models with dependent inputs. *Reliability Engineering & System Safety*, 107:115–121, 2012.

[O'Callaghan and Ramos, 2011] S. T O'Callaghan and F. T. Ramos. Continuous occupancy mapping with integral kernels. In *AAAI Conf. on Artificial Intelligence*, 2011.

[Rasmussen and Williams, 2006] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge, MA, 2006.

[Saltelli et al., 2000] A. Saltelli, K. Chan, M. E. Scott, et al. *Sensitivity analysis*, volume 134. Wiley New York, 2000.

[Schonlau and Welch, 2006] M. Schonlau and W. J. Welch. Screening the input variables to a computer model via analysis of variance and visualization. In *Screening*, pages 308–327. Springer New York, 2006.

[Sobol, 1990] I. M. Sobol. On sensitivity estimation for nonlinear mathematical models. *Matematicheskoe Modelirovanie*, 2(1), 1990.

[Vasudevan et al., 2009] S. Vasudevan, F. Ramos, E. Nettleton, and H. Durrant-Whyte. Gaussian process modeling of large-scale terrain. *J. Field Robotics*, 26(10):812–840, 2009.

[Xu and Gertner, 2008] C. Xu and G. Z. Gertner. Uncertainty and sensitivity analysis for models with correlated parameters. *Reliability Engineering & System Safety*, 93(10):1563–1573, 2008.